

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Context-Aware Systems and Methods, Location-Aware
Systems and Methods, Context-Aware Vehicles and
Methods of Operating the Same, and Location-Aware
Vehicles and Methods of Operating the Same**

Inventor(s):

Gopal Parupudi
Stephen S. Evans
Edward F. Reus

ATTORNEY'S DOCKET NO. MS1-698US

TECHNICAL FIELD

This invention relates to generally to the area of context-aware computing or ubiquitous computing.

BACKGROUND

The World Wide Web (WWW) was created to make content available from any source in any location around the world. Users of the Web are able to generally access a seemingly infinite number of resources via the Web. The Web has been highly successful in this regard. Yet, with the evolution of the Web, certain needs remain largely unmet. Specifically, people continue to have a need to access information that has a contextual aspect to it. That is, often times, individuals will find themselves in a computing environment that carries with it a certain context. Yet, the context of the environment cannot be easily incorporated into the present computing environment. As an example, consider the context of location. People generally have a need to access information, data, resources and the like, that have geographic dimensions to them. For example, individuals may desire to take advantage of services or products that are close in proximity to where they currently are located. In this regard, it is desirable to understand the individual's contextual location so that services, goods and the like can be made available to the individual. As "e-Commerce" continues to grow in importance, the necessity of bringing people, places, services and goods together in an efficient manner will become critically important.

To date, many attempts have been made to bring people, places, services and goods together. These various attempts have generally approached the problem from different directions in an often times incompatible manner. As an

example, consider the context of location. Some services have attempted to bring people and services together by defining large databases that maintain information about the services. For example, a list of restaurants may be maintained in a web accessible database where each restaurant is associated with a zip code in which the restaurant is located. When a user desires to locate a particular restaurant, they might simply enter the zip code where they are located to see a list of corresponding restaurants in that zip code. From the list of restaurants, they might be able to select one or two restaurants of interest. This approach is undesirable for a number of reasons. First, the operation of the system is dependent upon a central server that is responsible for receiving user queries and executing the queries to return the information to the user. In the event the server fails, so too does the service. In addition, this particular service might be suited to finding restaurants, but possibly not other businesses. In addition, the granularity with which the results are returned to the user may foist some of the search burden on the user (i.e. the user gets a list of restaurants in a nearby zip code, but has to further explore the list to select which ones are of interest). Further, the list of restaurants may include some restaurants that are blocked by some type of a physical barrier (i.e. a river, mountain, etc.) that makes the distance, as the crow flies, unroutable.

Providers of services and products want to be connected to nearby end-users. End-users want to consume these services and goods at the closest and most convenient location. Acquiring the services of a dentist or a plumber that lives somewhere "out on the net" is not appropriate if you need them to fill a cavity or unclog a sink. Looking for the nearest hotdog while in a stadium requires you to stay in the stadium.

1 There is an unsolved need to be able to create context-aware computing in
2 which computing devices can participate in their particular context. In specific
3 circumstances, there are needs to provide relational position awareness among
4 physical locations in both public and private views of the world. To date,
5 however, there is no one standardized view of the world that would unlock the
6 potential of context-aware computing. Context-aware computing is much more
7 than just position awareness—although this is a very big field in and of itself.

8 This invention arose out of concerns associated with developing a
9 standardized, context-aware infrastructure and related systems to unlock the
10 potential of context-aware computing.

11 12 SUMMARY

13 Context-aware systems and methods, location-aware systems and methods,
14 context-aware vehicles and methods of operating the same, and location-aware
15 vehicles and methods of operating the same are described. In various
16 embodiments, a context or location service module, implemented in software,
17 determines a vehicle context or a vehicle location based upon information that it
18 receives from various context providers or location providers respectively.
19 Software executing on a vehicle's computer can then cause one or more
20 applications that are associated with a vehicle computer to be modified in a
21 manner that changes their behavior. The behavior modification is based on the
22 current context or location of the vehicle and thus provides a context-specific or
23 location-specific user experience. The context or location can be ascertained
24 through the use of one or more hierarchical tree structures that comprises
25 individual nodes. Each node is associated with a context or location. The nodes

are traversable by the vehicle's software to ascertain a more complete context or location.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a diagram of an exemplary computing device that can be used in accordance with the described embodiments.

Fig. 2 is a conceptual diagram of an exemplary Master World and an exemplary Secondary World in accordance with the described embodiment.

Fig. 3 is an exemplary specific view of a Master World and a Secondary World and their relation to one another.

Fig. 4 is a flow diagram that describes steps in a method in accordance with the described embodiment.

Fig. 5 is a flow diagram that describes steps in a method in accordance with the described embodiment.

Fig. 6 is a high level diagram of an exemplary computing device architecture.

Fig. 7 is a somewhat more specific view of an exemplary computing device architecture.

Fig. 8 is a flow diagram that describes steps in a method in accordance with the described embodiment.

Fig. 9 is a flow diagram that describes steps in a method in accordance with the described embodiment.

Fig. 10 is a flow diagram that describes steps in a method in accordance with the described embodiment.

1 Fig. 11 is a side elevational view of an exemplary location beacon in
2 accordance with one embodiment.

3 Fig. 12 is a diagram of a vehicle computer in accordance with one
4 described embodiment.

5 Fig. 13 is a more detailed view of portions of the Fig. 12 vehicle computer.

6 Fig. 14 is a block diagram of an exemplary context-aware system in
7 accordance with one described embodiment.

8 Fig. 15 is a flow diagram that describes steps in a method in accordance
9 with one described embodiment.

10 Fig. 16 is a block diagram of an exemplary location-aware system in
11 accordance with one described embodiment.

12 Fig. 17 is a flow diagram that describes steps in a method in accordance
13 with one described embodiment.

14 Fig. 18 is a flow diagram that describes steps in a method in accordance
15 with one described embodiment.

16 Fig. 19 is a flow diagram that describes steps in a method in accordance
17 with one described embodiment.

18 Fig. 20 is a flow diagram that describes steps in a method in accordance
19 with one described embodiment.

20 Fig. 21 is a flow diagram that describes steps in a method in accordance
21 with one described embodiment.

22 Fig. 22 is a diagram of an exemplary system in accordance with one
23 described embodiment.
24
25

DETAILED DESCRIPTION

Overview

To provide a standardized solution, embodiments described just below provide a uniform definition of the world. The uniform definition is defined in terms of a hierarchical tree of nodes, where each node represents some aspect of the world. Each node is connected to at least one other node by a branch. An exemplary classification of nodes takes place on a physical level (e.g. physical locations such as political entities, infrastructure entities and public places), as well as a non-physical level (e.g. military APOs). This hierarchical nodal structure is referred to as the Master World, and is a standardized view worldwide. Each node of the Master World has various attributes associated with it that assist in context-aware computing. Exemplary attributes include a unique ID, name, geographic entity class, latitude/longitude, relative importance, contextual parents to name just a few. The Master World is useful because it can be used to determine the relative location of a place anywhere in the world and at any definable granularity.

Once an individual's location or a place an individual is interested in is determined, various services that reference the location can be offered to the individual based on their location. That is, value is provided by the Master World model in the ability to tie services to nodal locations in the Master World.

Building on this concept, two additional concepts add value—the concept of so-called Secondary Worlds and a “geozone.”

A Secondary World is a powerful computing mechanism whereby individual entities (such as businesses or organizations) can define their own particular worlds that need not necessarily conform to the Master World view of

1 the world. That is, while the Master World is essentially a physical hierarchical
2 representation of the world, the Secondary Worlds can be physical and/or logical
3 representations of each individual entities' world view. One particularly useful
4 aspect of the Secondary World is that it links, at at least one point, into the Master
5 World. Thus, within any Secondary World, a user's location not only within the
6 Secondary World, but the Master World as well can be determined. Various
7 services can be attached to the nodes of the Secondary World. Based upon a user's
8 calculated position, these various services that are associated with Secondary
9 World nodes can be offered to the user. In addition, because the user's context is
10 determined relative to the Master World, other services that may not be associated
11 with a particular Secondary World can be offered.

12 A geozone is essentially a spatial indexing mechanism by which the Master
13 World is subdivided into individual zones. In the described embodiment, the
14 zones are subdivided through the use of a quadtree algorithm that is dependent on
15 a density function (although many other spatial index approaches can also be
16 used). Once a desired density level is achieved (density might be defined in terms
17 of points of interest per zone), each node on the Master World is assigned a
18 particular geozone. Geozones enable proximity calculations to be computed in a
19 fast and straight forward manner.

20 A useful aspect of the Master and Secondary Worlds are that they are
21 "reachable" from various computing devices such as stationary (i.e. desktop
22 devices) or mobile computing devices (i.e. cell phones, laptops etc.). That is, the
23 Master World (or at least a portion of it) and one or more Secondary Worlds can be
24 either locally maintained on the computing device, or accessed, e.g. via the Web or
25 some other mechanism, so that a user can derive their context. For example, the

1 Secondary World can be downloaded onto the computing device so that a user can
2 derive their context within the Secondary World. Once a user's context is
3 determined from the Master World and one or more Secondary Worlds, a various
4 robust collection of context-aware solutions become available to the user. For
5 example, specific Secondary World services can be offered or Master World
6 services can be offered. Additionally, services from other Secondary Worlds might
7 also be offered since the user's location may be known (or made known) to these
8 other Secondary Worlds. In this way, the Master World can link two or more
9 Secondary Worlds together.

10 Another aspect is that the described embodiments harness the computing
11 power of each computing device in determining the device's location. Here, by
12 virtue of having the Master World and one or more Secondary Worlds reachable
13 by the device (and possibly locally maintained on the device), the device itself
14 determines its own context.

15 One embodiment provides a client side device that is configured to utilize
16 the context-aware structures that are discussed above, i.e. the Master and one or
17 more Secondary Worlds. The Master World or a portion thereof can be locally
18 available on the device or can be accessible at another location, e.g. via the Web.
19 In this embodiment, the client device has a location service embodied thereon.
20 The described location service is a software module that can determine the
21 location of the device and can answer queries from various applications (either
22 executing on the device or off the device). The location service determines the
23 location of the device by using the Master World and one or more Secondary
24 Worlds. The applications query the location service through one or more
25

1 Application Program Interfaces (APIs) or Events to get location information that is
2 used by the applications to render a service.

3 The location service makes use of one or more location providers that
4 convey information to the device. This information can be information that is
5 specific to the location provider, or can be information that can be mapped directly
6 into a node of the Master World or Secondary Worlds. Exemplary location
7 providers can include Global Positioning Service (GPS) providers, cell phone
8 providers (cell providers), Bluetooth providers, a user interface provider and the
9 like. The location providers provide information that gives some aspect of a
10 device's current location. This information is used by the location service to
11 ascertain the location of the device.

12 One particularly advantageous feature of the client device is a standard or
13 common location provider interface. The location provider interface enables the
14 various location providers to provide information to the location service so that the
15 location service can use the information to determine its location. Essentially, the
16 multiple location provider interface is a common interface that enables multiple
17 different location providers to provide location information (or hints) about
18 location to a location service that is on a device. The location providers can
19 provide the location information constantly, at intervals, or when polled by the
20 device. The location information can be provided with confidence and accuracy
21 estimates to enable the location service to evaluate the relative quality of the
22 information before it is used. The various providers also have the ability to self-
23 monitor themselves which assists in the providers' ability to intelligently convey
24 information to the location service. By having a common interface, the collection
25 of location providers is dynamically extensible—that is location providers can be

1 added or removed from the collection of location providers without any
2 interference of the functionality performed by the location service or device. The
3 location providers can be added or removed while the device is operating. This is
4 particularly useful in accommodating location providers that are developed in the
5 future. In this particular embodiment, two levels of abstraction are provided i.e.
6 (1) the provider interface that receives information from the location providers and
7 (2) the API/events layer that enables applications to get at the various information.

8 One focus of this embodiment is a device that can collect context
9 information (e.g. location information) from a variety of different sources,
10 determine the device's current context from that information, and provide the
11 current context at some level to one or more applications that can use the device's
12 context to render a service or enable the device to participate in its context
13 environment.

14 In the described embodiment, the device receives location information or
15 hints about its location. This information is collated and mapped by the location
16 service into a node in the Master World and/or Secondary World. The hierarchical
17 trees can then be traversed to determine the device's accurate location in both the
18 Secondary World and the Master World. At this point, the device has determined
19 its context. The information that is collected can be subject to arbitration to ensure
20 that only highly trusted information is used to determine context. The location
21 information can be cached to provide "current location information" which, for a
22 definable period of time will be accurate to some degree. Thus, if for some reason
23 other location providers are unavailable, the cache can be used to ascertain
24 location.
25

1 Once a device's location is determined, the device can apply a security
2 policy to the information. Once this is done, the device can begin to answer
3 queries from various applications.

4 One aspect of the described embodiment is a "favorite locations" aspect in
5 which the device can be automatically configured, when it determines its context,
6 so that it can adjust to the different locations.

7 Further, various types of location providers can convey different types of
8 information. For example, a so-called "thin provider" provides location
9 information that is translated by the location service into the appropriate node
10 information. A so-called "thick provider" includes logic that takes location
11 information and provides it in a form that can map directly into the Master World
12 or Secondary World.

13 In another embodiment location translation services are provided that are
14 directed to determining, as accurately as possible, the context or location of the
15 device. In this embodiment, information is received from the various location
16 providers. This information includes location, accuracy and confidence (all of
17 which are provided by the location provider), trust (which is assigned to a location
18 provider by the device or a user) and a timestamp (which helps to age the location
19 information). The location translation processing involves determining which of
20 the location providers are valid and active. The location providers can be ranked
21 in accordance with the confidence and trust levels. This defines an ordered list of
22 location providers. Provision is made for a situation in which all of the location
23 providers may go inactive. If so, a "current location" is used as a location
24 provider whose confidence decreases over time.

1 In the event that information from two or more of the location providers
2 conflicts, then measures can be taken to use information for which there is a
3 higher level of trust. The information that is provided by all of the location
4 providers (assuming no conflict) can then be used to determine a tree structure and
5 a node's entity ID (EID). The tree might be the Master World and the EID is a
6 node on the Master World. The tree might also be a Secondary World and the EID
7 (or location unique identifier or "LUID") is a node on the Secondary World.
8 Once this information is collected, complete location information can be
9 determined by simply traversing the tree(s). Once a device's location is
10 determined, a cache can be updated with the current location (including a time
11 stamp).

12 In another embodiment, privacy issues in the context-aware computing
13 environment are addressed. In this embodiment, the location service has acquired
14 location information that pertains to the location of a particular device. A privacy
15 manager determines what level of information to provide to applications that
16 might request the information. The privacy manager can reside on the computing
17 device itself, or can be proxied by a trusted third party.

18 According to this embodiment, a scale of privacy levels are defined. Each
19 level is defined to include more or less specific information about the location of a
20 particular device. A user is able to assign a privacy level to entities that might
21 request location information. Additionally, each node of the Master World and a
22 Secondary World can have a privacy level associated with it. When a query from
23 an application is received, the privacy manager first determines who the query is
24 from and the privacy level associated with the application or entity. The privacy
25 manager then evaluates one or more of the Master World and the Secondary World

1 to find a node that has a corresponding privacy level. When a corresponding node
2 is found, information at that particular granularity is provided to the requesting
3 application or entity.

4 In another embodiment systems and methods of providing a location
5 provider in the form of a location beacon are described. In this embodiment, a
6 location beacon is provided that can be mounted in various areas (public/private
7 areas) to beacon the location to any computing devices within transmission range.
8 The information that is transmitted enables a device to determine its location or
9 context. The location beacon can transmit information that is specific to the
10 location service that uses the information. Transmitted information can include an
11 EID/URL pair, and a LUID/URL pair. The EID gives the node identification of a
12 node in the Master World; and, the associated URL gives a protocol to
13 communicate with the Master World. The URL might, for instance, link to a
14 server that can provide additional context information that uses the EID. The
15 LUID indicates a node on a Secondary World that corresponds to a current
16 location; and the URL gives a protocol to communicate with the Secondary World.
17 For example, the URL can link with a server that is hosting the Secondary World.
18 This server can then be queried to discover more information about the Secondary
19 World (i.e. Secondary World tree structure, location of associated resources, etc.)
20 With the EID and LUID (along with the URLs), a device can now traverse the
21 Master World or Secondary World to determine its location. Various technologies
22 can be used to implement the beacon (wireless, RF, IR). The beacon can be a
23 "program once" device to deter tampering. Programmable beacons can, however,
24 be provided. Security can also be provided in the form of a verifiable signature
25

1 that is provided with the beacon information to assure the veracity of the
2 transmitted information.

3 A useful context-aware computing aspect of the beacon is the concept of
4 "location-enabled access". That is, in addition to (or separately from) receiving
5 location information, a beacon can transmit code download pointers that enable
6 smart devices to access software code that allows the device to participate in its
7 current context.

8 9 **Exemplary Computing System**

10 In the context of this document, the term "computing device" is used to
11 refer generally to any type of computing device. Characteristics of exemplary
12 computing devices are that they typically include one or more processors,
13 computer-readable media (such as storage devices and memory), and software
14 executing on the one or more processors that cause the processors to implement a
15 programmed functionality. In particular embodiments, implementation takes place
16 in the context of mobile computing devices (e.g. laptop computers and the like),
17 and/or hand-held computing devices (e.g. palm PCs, wireless telephones and the
18 like).

19 Fig. 1 is a schematic diagram that constitutes but one example of a
20 computing device that is suitable for use in connection with the described
21 embodiments. It is to be understood that portions of the illustrated computing
22 device can be incorporated in one or more of the computing devices (e.g. palm
23 PCs, wireless telephones, etc.) with which particular embodiments are envisioned
24 for use.
25

1 Computer 130 includes one or more processors or processing units 132, a
2 system memory 134, and a bus 136 that couples various system components
3 including the system memory 134 to processors 132. The bus 136 represents one
4 or more of any of several types of bus structures, including a memory bus or
5 memory controller, a peripheral bus, an accelerated graphics port, and a processor
6 or local bus using any of a variety of bus architectures. The system memory 134
7 includes read only memory (ROM) 138 and random access memory (RAM) 140.
8 A basic input/output system (BIOS) 142, containing the basic routines that help to
9 transfer information between elements within computer 130, such as during start-
10 up, is stored in ROM 138.

11 Computer 130 further includes a hard disk drive 144 for reading from and
12 writing to a hard disk (not shown), a magnetic disk drive 146 for reading from and
13 writing to a removable magnetic disk 148, and an optical disk drive 150 for
14 reading from or writing to a removable optical disk 152 such as a CD ROM or
15 other optical media. The hard disk drive 144, magnetic disk drive 146, and optical
16 disk drive 150 are connected to the bus 136 by an SCSI interface 154 or some
17 other appropriate interface. The drives and their associated computer-readable
18 media provide nonvolatile storage of computer-readable instructions, data
19 structures, program modules and other data for computer 130. Although the
20 exemplary environment described herein employs a hard disk, a removable
21 magnetic disk 148 and a removable optical disk 152, it should be appreciated by
22 those skilled in the art that other types of computer-readable media which can
23 store data that is accessible by a computer, such as magnetic cassettes, flash
24 memory cards, digital video disks, random access memories (RAMs), read only
25

1 memories (ROMs), and the like, may also be used in the exemplary operating
2 environment.

3 A number of program modules may be stored on the hard disk 144,
4 magnetic disk 148, optical disk 152, ROM 138, or RAM 140, including an
5 operating system 158, one or more application programs 160, other program
6 modules 162, and program data 164. A user may enter commands and information
7 into computer 130 through input devices such as a keyboard 166 and a pointing
8 device 168. Other input devices (not shown) may include a microphone, joystick,
9 game pad, satellite dish, scanner, or the like. These and other input devices are
10 connected to the processing unit 132 through an interface 170 that is coupled to
11 the bus 136. A monitor 172 or other type of display device is also connected to the
12 bus 136 via an interface, such as a video adapter 174. In addition to the monitor,
13 personal computers typically include other peripheral output devices (not shown)
14 such as speakers and printers.

15 Computer 130 commonly operates in a networked environment using
16 logical connections to one or more remote computers, such as a remote computer
17 176. The remote computer 176 may be another personal computer, a server, a
18 router, a network PC, a peer device or other common network node, and typically
19 includes many or all of the elements described above relative to computer 130,
20 although only a memory storage device 178 has been illustrated in Fig. 1. The
21 logical connections depicted in Fig. 1 include a local area network (LAN) 180 and
22 a wide area network (WAN) 182. Such networking environments are
23 commonplace in offices, enterprise-wide computer networks, intranets, and the
24 Internet.

1 When used in a LAN networking environment, computer 130 is connected
2 to the local network 180 through a network interface or adapter 184. When used
3 in a WAN networking environment, computer 130 typically includes a modem 186
4 or other means for establishing communications over the wide area network 182,
5 such as the Internet. The modem 186, which may be internal or external, is
6 connected to the bus 136 via a serial port interface 156. In a networked
7 environment, program modules depicted relative to the personal computer 130, or
8 portions thereof, may be stored in the remote memory storage device. It will be
9 appreciated that the network connections shown are exemplary and other means of
10 establishing a communications link between the computers may be used.

11 Generally, the data processors of computer 130 are programmed by means
12 of instructions stored at different times in the various computer-readable storage
13 media of the computer. Programs and operating systems are typically distributed,
14 for example, on floppy disks or CD-ROMs. From there, they are installed or
15 loaded into the secondary memory of a computer. At execution, they are loaded at
16 least partially into the computer's primary electronic memory. The invention
17 described herein includes these and other various types of computer-readable
18 storage media when such media contain instructions or programs for implementing
19 the steps described below in conjunction with a microprocessor or other data
20 processor. The invention also includes the computer itself when programmed
21 according to the methods and techniques described below.

22 For purposes of illustration, programs and other executable program
23 components such as the operating system are illustrated herein as discrete blocks,
24 although it is recognized that such programs and components reside at various
25

1 times in different storage components of the computer, and are executed by the
2 data processor(s) of the computer.

3 4 **Defining the World**

5 One of the problems to date with attempting to solve the context-aware
6 computing problem is that every proposed solution has its own approach, data
7 structures, processes and the like. There is little if any standardization between the
8 various approaches. In the described embodiment, standardization is achieved at
9 the foundational level by defining a universal view of the Earth. That is, a
10 universally acceptable definition of the Earth is proposed and is useable in various
11 computing scenarios to enable context-dependent computing. In this document, a
12 specific example of context-dependent computing is given in the form of location
13 dependent computing. It is to be understood that this constitutes but one example
14 of a context in which the various embodiments discussed below can be employed.
15 Other “contexts” can include, any information that can fit into a hierarchical
16 structure including, without limitation, role/personnel in an organization, device
17 categorizations, current activity, current environment, active devices and the like.

18 19 **The Master World**

20 A Master World is defined as a politically correct and publicly accepted
21 hierarchical tree structure that catalogs physical location or geographic divisions
22 of the Earth. The Master World is defined in such a way that many different
23 classes of political, administrative and geographic entities across the entire Earth
24 are included. Areas of political contention are accounted for by presenting a view
25 of the world based on the language/locale of the computing device.

Fig. 2 shows an exemplary hierarchical tree structure 200 that represents a portion of the Master World. The Master World contains multiple nodes 202, with each node representing some type of geographic division (e.g. political or natural entity) of the Earth. In the illustrated example, the nodes of the Master World are arranged in the following groups: (1) political or natural entities (e.g. continents, countries, oceans, states, counties, cities and the like); (2) infrastructure entities (e.g. postal codes, area codes, time zones and the like); (3) public place entities (e.g. parks, malls, airports, stadiums, and the like); and (4) non-physical entities (military postal code regions, vacation regions, affiliate coverage areas of television networks that can be geographically discontinuous, and the like).

In the Fig. 2 example, the top node of the tree structure represents the Earth. Each node underneath the top node represents a geographical division of the Earth. In this example, none of the nodes have an association with any businesses or services. That is, there is a distinction between node entities that are part of the Master World and non-geographic places where activities take place. Though the Master World includes nodes for public places (i.e. airports, malls, etc), it does not include individual listings of businesses or service providers. Each node is uniquely identified by an ID (EID or entity ID). In addition to the unique EIDs, a URL is associated with the tree structure and provides a context for the tree structure as will become apparent below.

As an example, consider the following: Seattle-Tacoma International Airport (SeaTac) will be included in the Master World, but references to individual airline business locations at SeaTac might be “leaves” on the tree that are tagged by the SeaTac Airport EID (see “Secondary World” section and the Table below). Similarly, the Seattle Center might be a node on the Master World, while the

1 Seattle Arts Festival, Bumpershoot, the Seattle Sonics NBA Team, and the Seattle
2 Center Starbucks Coffee Shop might be tagged with the Seattle Center EID. As
3 another example, the Master World also contains nodes for all Interstate
4 (motorway) exits. For example, the I-90, Exit 109, Washington is a node in the
5 Master World. The Best Western Inn located at 1700 Canyon Road in Ellensburg,
6 Washington might be tagged with the EID of this Exit.

7 Thus, the Master World provides a uniform way of defining locations. The
8 uniform location definitions can then be universally used to assign attributes to
9 goods or services. Whenever a computing device determines its location to
10 correspond to a particular uniform location definition, it can take advantage of the
11 location-dependent goods or services that share the uniform location definition.
12 The Master World is useful because it is a standardized view of the world. Its
13 accurate standardized geographic dimension attribution can be easily accessed by
14 both providers and consumers. Services and product providers (or third parties
15 such as search engines, network and yellow-page database directories) can use the
16 nodes of the Master World by assigning a standardized persistent geographic
17 reference to all commerce locations or points of interest. These commerce
18 locations or points of interest can be considered as “leaves” on the tree structure.

19 In the illustrated example, the nodes of the Master World have one or more
20 attributes that facilitate its use. Exemplary attributes are described in the table
21 immediately below:

Attribute	Description
Entity ID (EID)	The EID is a unique ID for each node. No two nodes have the same EID.
Name	The name is defined in terms of the neutral ground truth (NGT) name. The NGT name supports various language translations for entity names as appropriate (e.g. Pacific Ocean, Pazifischer Ozean, Oceano Pacifico, etc.)

Geographical Entity Class (GEC)	The GEC is a geographical classification of each node. An exemplary GEC is discussed below in the "Geozone" section.
Latitude	The horizontal coordinate position on the globe (i.e. the coordinate position of the node's centroid)
Longitude	The vertical coordinate position on the globe (i.e. the coordinate position of the node's centroid)
Relative Importance	The geographic importance of an entity in reference to other entities in the same region. Value from 1 to 256 (e.g. New York City = 3, Los Angeles = 4, and Omaha = 5 even though Omaha is much smaller but almost as important in relation to surrounding populated places)
Contextual Parent(s)	The parents of the parent/child relationship for each node. Multiple parents are supported (e.g. Redmond is a child of King County, Area Code 425, the Pacific Time Zone, and the MSNBC affiliate KING TV).
Source	The source of origin for the record (e.g. Microsoft or a specified data vendor)
Start Date	Date when the node information was first valid
End Date	Date when the node information was last valid (retired zip codes, breakup of countries)
Modification Date	Records date changes that are made to the record relating to retirement or updates to any fields
Status	Active, lashed (links duplicate nodes together), pending or retired

The attributes listed above constitute exemplary attributes only. Other attributes that are different from and/or additional to those referenced above could be used. A few exemplary entity or node records that employ the above attributes are shown below:

Entity ID (EID)	24948
Name	Pacific Ocean, Pazifischer Ozean, Oceano Pacifico, etc.
Geographical Entity Class (GEC)	138/Ocean
Latitude	0 (+000° 00' 00")
Longitude	-170 (-170° 00' 00")
Relative Importance	1
Contextual Parent(s)	World
Source	MSFT GeoUnit
Start Date	0/0/00
End Date	0/0/00
Modification Date	01/18/00
Status	Active

Entity (EID)	ID	27490
Name		Redmond
Geographical Entity Class (GEC)		78/non-capital town
Latitude		47.6768303 (+047° 40' 36")
Longitude		-122.1099625 (-122° 06' 35")
Relative Importance		107
Contextual Parent(s)		1. King, second level [Washington, United States] 2. Puget Sound-Seattle, travel region [Washington, United States]
Source		MSFT GeoUnit
Start Date		0/0/00
End Date		0/0/00
Modification Date		01/18/00
Status		Active

The Master World also serves as a repository of common denominator links between itself and various "Secondary Worlds" and as a conduit that connects Secondary Worlds to other Secondary Worlds. Content, service and device providers can use the Master World to associate their publicly available offerings with a geographic location and the corresponding multiple branch hierarchical structure. This location will be associated with a single entity within the tree structure thereby allowing geographic and time/distance calculations and the necessary parent/child relationship navigation.

The Master World Index (Geozones)

By definition, the Master World provides a hierarchical structure of entities (nodes) that cover the entire globe. Upward navigation within the hierarchy is

1 quite natural. Efficient navigation downward requires geographic proximity
2 awareness. Additionally, there are possible scenarios that will require jumping
3 from branch to branch in order to successfully return values in a query, or for more
4 accurate calculations of distances to close "leaves" attached to nodes other than
5 the original source node. The Master World makes use of an index scheme that
6 can identify peer level nodes by virtue of the geographical proximity. This
7 indexing scheme makes use of a quad tree algorithm to define so-called
8 "geozones."

9 A quadtree is essentially a spatial index that breaks coverage into
10 homogeneous cells of regularly decreasing size. Each quadrant of the tree has up
11 to four children. The quadtree segmentation process can continue until the entire
12 map is partitioned based on many different end result criteria including the density
13 of the number of items (e.g. points of interest) in each quad. The approach
14 provides a form of spatial index that accelerates spatial selection and content
15 identification.

16 To complete the spatial indexing scheme to provide each node with a
17 defined geozone, a quadtree algorithm is applied to the nodes and can be based
18 upon a desired density of, for example, points of interest that are to occur in any
19 one zone. Once all of the zones have been defined, each zone is given a unique ID
20 (e.g. top/left and bottom/right Latitude and Longitude pairs). Each of the nodes of
21 the Master World is then assigned a zone in which it is located. Quadtree
22 algorithms are known and will be appreciated by those of skill in the art.

The Master World Location

As can be appreciated, having a uniform standardized representation of the world in the form of a hierarchical traversable tree structure can greatly facilitate the manner to which context-dependent, and more specifically, location-dependent goods and services can be linked.

In the described embodiment, a computing device has access to at least a portion of the Master World. For example, the computing device can have the Master World saved in an internal storage device, it can comprise part of the computing device's operating system, or the device might access the Master World via a network medium such as the Internet. With the Master World tree structure being accessible to each computing device, each device has the power to determine its own context or node-referenced location. That is, the computing device can determine, through software it is executing, its particular location, i.e. node. Once the computing device determines an associated node, it can simply traverse the tree to ascertain its complete location.

For example, if a computing device determines that it is currently located at a node that corresponds to the City of Redmond, it can traverse the Master World tree structure to ascertain that it is in the State of Washington, Country of The United States, on the continent of North America. By ascertaining its precise location, the computing device (or its user) is now in a position to take advantage of location-dependent services that might be offered. This particular model is a tremendous improvement over current models that utilize a central server to ascertain location for a number of different devices. In that model, each device (or user) provides information about its location (e.g. perhaps the user enters the zip code or city that the device is currently in) and might enter a query to find, for

1 example, a McDonald's restaurant in his zip code. The server then takes this
2 information and might, for example, tell the user about the location of all of the
3 McDonald's restaurants within that zip code or city. If the servers fails in this
4 model, then none of the computing devices can take advantage of its services. In
5 the present model, each computing device is self-sustaining. Each can determine
6 its own location, and accordingly, each device can take advantage of location-
7 dependent services. For example, if the computing device understands that it is
8 located on a particular node of the Master World, then it can execute queries to
9 find a McDonald's that has an EID that corresponds to the particular node in
10 which the computing device is located. Particular robustness is provided through
11 the use of the above-described geo-zones. The geo-zones enable proximate
12 geographic divisions to be quickly searched in an efficient manner. For example,
13 if an individual is looking for the nearest Kinko's to make copies and none are
14 located in the geo-zone that corresponds to the node in which the computing
15 device is located, then adjacent geo-zones can be quickly searched.

17 Secondary Worlds

18 In the described embodiment, the concept of a Secondary World is used to
19 provide support for additional context. A secondary world might be defined by a
20 third party organization or company and contains nodes that comprise physical
21 and/or logical entities that are unique to that organization. The nodes of the
22 Secondary World may or may not have much context outside of the particular
23 organization that defined the Secondary World, since a secondary world could be
24 made either public or private. The Secondary Worlds do not duplicate the Master
25 World, but rather supplement it in a unique, organization specific manner. While

1 the Master World is defined to be a widely accepted standard, each Secondary
2 World can be a widely variant representation of an organization's proprietary view
3 of the world. In the described embodiment, each Secondary World has at least one
4 node that is linked with a node of the Master World. This gives the Secondary
5 World a context or location in the Master World. Also note that in some context
6 applications, several secondary worlds may be accessed, each providing additional
7 context specific pieces of location data.

8 Fig. 2 shows an exemplary Secondary World 204 that comprises a plurality
9 of nodes 206. Each of the nodes 206 constitutes a physical or logical entity. For
10 example, the nodes can constitute a company, its divisions, regions campuses,
11 buildings, floors in various buildings and rooms on various floors. At least one of
12 the nodes is linked with a node of the Master World. The nodes of the Secondary
13 World can have the same attributes as the nodes of the Master World.

14 As an example of a Secondary World, consider that Boeing might define a
15 Secondary World that includes a list of entities that are important to its employees.
16 The root entity would be "Boeing Corp." and its children might be company
17 divisions (St. Louis Military Division, Everett Plant, Corporate HQ, etc.). Further
18 down the tree structure, individual nodes might be defined to represent individual
19 buildings (Hanger 12), offices within this building (Office 1001), building areas
20 (Southwestern quadrant of hanger 12), etc. Each entity or node has a unique
21 identifier (Local Unique ID or "LUID") and a URL that is associated with the tree
22 on which the node occurs. The URL uniquely identifies the Secondary World tree
23 structure so that a user within that world can determine how to interact with the
24 world. This aspect is discussed below in more detail. Boeing can then use the
25

1 LUIDs to associate equipment, services, departments or even personnel to a
2 physical or logical location.

3 As a more concrete example, consider Fig. 3 which shows an exemplary
4 portion of the Master World 300 and a Secondary World 302. Master World 300
5 includes the following nodes: World, United States, Washington, Redmond, and
6 Zip = 98052. The exemplary Secondary World 302 is a hierarchical tree structure
7 that has been defined by Microsoft Corporation and includes the following nodes:
8 Microsoft, Redmond Campus, 1 Microsoft Way, Building 26, 3rd floor, Conference
9 Room 3173, Building 24, 2nd floor, Conference Room 1342. In this example, the
10 Secondary World 302 “touch points” into the Master World from the Redmond
11 node. In this example, a video projector is shown as being associated with the
12 node “Conference room 1342”. Here, the video projector is not a node in the
13 secondary world. Rather, the video projector is an item in some other resource
14 discovery service (e.g. the active directory) and includes a location attribute that is
15 a pointer to “Conference room 1342.” There may be times, however, when nodes
16 can be created in the worlds to represent the location of key services—the node
17 themselves, however, would not represent the services.

18 Like the Master World, the Secondary World is advantageously accessible
19 to a user’s computing device. It could, for example, be downloaded—completely
20 or partially--and stored on a storage device and accessed when needed. It might
21 be downloaded for a one time use only. The Secondary World enables the
22 computing device to ascertain its context within the Secondary World. In this
23 example, the computing device would, by using the Secondary World, compute its
24 location within the Secondary World. The computing device can do this by
25 traversing the tree structure from the node in which it is currently located to the

1 root node. This would, for example, give the computing device (and hence the
2 user) a complete Secondary World context. Once the Secondary World location is
3 known, the user is in a position to take advantage of goods or services that are
4 associated with the nodes of the Secondary World. That is, once the computing
5 device determines its Secondary World context, it is ready to become an active
6 participant in the Secondary World.

7 Tremendous value can be achieved by associating goods or services with
8 the individual nodes of the Secondary World. For example, Conference Room
9 1342 has a video projector associated with it. That is, the location of the video
10 projector is in Conference Room 1342. Assume that an individual in Conference
11 Room 3173 has a presentation that requires the use of the video projector such as
12 the one located in Conference Room 1342. Normally, an individual would have
13 no way of ascertaining the location of the video projector other than perhaps
14 physically calling over to the building to check whether there is a video projector
15 available. In this example, because the user's computing device is able to
16 ascertain its location within the Secondary World, it is able to locate the video
17 projector in Conference Room 1342. It would do this by simply executing
18 software that traverses the Secondary World tree structure to find the resource of
19 interest.

20 Note also that because there is a link into the Master World, the computing
21 device is able to derive its context (location) within both worlds. This enables the
22 computing device, and hence the user, to take advantage of goods and services that
23 are associated with the Secondary World, as well as participate in location-
24 dependent services that are consumable based upon the user's location in the
25 Master World.

1 Fig. 4 is a flow diagram that describes steps in a method in accordance with
2 the described embodiment. The steps described just below are implemented by a
3 computing device which, in the illustrated example, is a hand-held mobile
4 computing device.

5 Step 400 accesses first and second hierarchical tree structures that are
6 resident on a computer-readable media. In this example, the tree structures might
7 be stored on the device or might be accessible via a network such as the Internet.
8 An exemplary first tree structure is the Master World and an exemplary second
9 tree structure is a Secondary World. Alternately, the tree structures could both be
10 Secondary Worlds. Once the tree structures have been accessed by the device,
11 step 402 traverses multiple nodes of the tree structures to derive the context of the
12 computing device. In this example, the computing device receives information
13 that informs it as to its location at a node of one of the trees. This information can
14 come to the computing device in any suitable way, e.g. a user can enter the
15 information through a User Interface (UI) or the location might be broadcast to the
16 computing device by another computing device (e.g. through the use of Bluetooth
17 technology or Universal Plug and Play (UpnP). Specific examples of how this
18 information can be conveyed to the computing device are given below in more
19 detail. Regardless of how this information is conveyed to the computing device,
20 once the computing device has the information, it executes software that traverses
21 one or both of the tree structures to derive its context which, in this example, is the
22 device's location.

Defining Secondary Worlds

As was mentioned above, one particularly valuable aspect of the described embodiment is that individual organizations can define their own Secondary Worlds. This gives the organization a great deal of flexibility in providing goods and services and, more broadly, increasing the efficiency of their organization. In one embodiment, a software tool is provided that enables individual organizations to define and maintain their own Secondary Worlds.

In one embodiment, each secondary world can be uniquely identified as a name space (e.g. an XML namespace). This ensures that any overlap in names between the Secondary World and the Master World will not result in a collision. As an example, consider the following: the Master World might contain an entity identified as “Chicago” referring the city. A Secondary World that is established by the National Basketball Association (NBA) and a different Secondary World that is established by the Caterpillar Corporation might also have entities named “Chicago” that refer to completely different entities than the Master World’s “Chicago.” For example, the NBA’s “Chicago” might refer to an NBA market area while Caterpillar’s “Chicago” might refer to a sales district. Having the namespace separation between the Master and Secondary Worlds can ensure that there not a collision between identically named entities because each name space is uniquely different from every other namespace.

Fig. 5 is a flow diagram that describes steps in a method of building a context-aware data structure. These steps are implemented by a software tool that is executing on a computing device.

Step 500 receives input from a source that specifies information that pertains to physical and/or logical entities. In this example, a system administrator

1 might physically enter information about the structure of the Secondary World that
2 they desire to define. This information can include information about buildings,
3 divisions, conference rooms and the like. Step 502 then processes the information
4 to define a hierarchical tree structure that has a context. In this example, the
5 context is location. It will be appreciated, however, that other contexts could be
6 employed. Each of the nodes in the hierarchical tree structure represents a
7 separate physical or logical entity. Step 504 then links at least one of the nodes of
8 the hierarchical tree structure with another tree structure having a context. In this
9 example, this other tree structure can comprise the Master World. Once the tree
10 structures have been built and linked, they are ready for traversal in a manner that
11 enables context to be derived from one or more of the nodes.

12 **Location as a Service**

13
14 In the above examples, the computing device is able to determine its own
15 location. In the embodiment about to be described, the computing device
16 determines its location by using location information that is provided to it from a
17 number of different sources of information. The device is able to take the
18 information that is provided to it and process the information to determine a
19 particular node on one or more hierarchical trees. Once the device has done this, it
20 can determine its complete location which is a useful thing to know particularly
21 when there are location-dependent services that can be consumed by the device's
22 user.

23 Fig. 6 shows a high level diagram of an exemplary computing device 600
24 that comprises, among other components, a context service module 602 and one or
25 more context providers 604. The context service module 602 can be implemented

1 in any suitable hardware, software, firmware or combination thereof. In this
2 particular example, the context service module is implemented in software that is
3 executed by one or more device processors. The context service module 602
4 receives context information from one or more context providers 604 and
5 processes the information to determine a current device context. In this particular
6 example, the device context is the device's location. Accordingly, the context
7 providers are location providers that provide location information, in various
8 forms, to the context service module 602 for processing. The location providers
9 604 receive information from various sources of context information (location
10 information) 606.

11 In the context of this document, a context provider comprises a software
12 component that can either be implemented on the device or off the device. The
13 context provider can also include any suitable hardware, firmware or combination
14 thereof. The role of the context providers are to receive information from sources
15 606 and convey the information to the context service module 602 so that the
16 context service module can use the information to determine a current device
17 context.

18 In the case where the context of the device is the device's location, sources
19 606 provide various information to the location providers 604 that pertains to the
20 device's current location. As an example, the sources of the information can
21 include various information transmitters such as a GPS system, cell phone or cell
22 ID, wireless transmitters that transmit location information, location beacons,
23 802.11 transmitters and various other sources of information. The sources of
24 information can also include other computing devices that might, for example,
25 provide location information through Bluetooth technology. In addition, a source

1 of information 606 might include a person who, for example, physically enters
2 location information into the device 600 so that the device can process the
3 information to determine its location.

4 When the device 600 receives the location information from the sources
5 606, it processes the information with the location providers 604 and provides the
6 information to the location service module 602. The location service module 602
7 processes the location information and determines a particular node on one or
8 more of the hierarchical tree structures to which it has access which corresponds to
9 its current location. The location service module 602 can then traverse the tree
10 structures to determine a complete location for the device. Once the complete
11 location is determined, the device 600 can begin to interact with one or more
12 applications 608 that can query the device about its particular location so that one
13 or more location-dependent services can be rendered to the device. In this
14 example, the applications 608 are illustrated as being separate from the device. It
15 is to be understood, however, that the applications could be executing on the
16 device, e.g. a browser application.

17 As shown, the applications 608 can make calls to the device to ask the
18 device where it is located. The device is configured to receive the calls and
19 respond in an appropriate manner to the application. Once the application has the
20 device's location information, it can then render location specific services to the
21 device.

22 **Consider the following example:** You are a traveler and have a hand-held
23 mobile computing device that contains a Master World tree and a Secondary
24 World tree for SeaTac International Airport. You are scheduled to depart on a
25 plane for China from Concourse C. SeaTac International Airport has designed its

1 Secondary World to have the following nodes: "Arrivals", "Departures",
2 "Concourses", "Airlines", "Gates assigned to Airlines", and "Gate Location".

3 When you arrive at the airport, as you enter the airport your computing device
4 receives location information from different sources and with that information
5 your device determines that your location is in the Arrivals node. SeaTac
6 International has bank of servers that are executing applications to assist you while
7 you are in the airport. There are applications that can help you find services,
8 locate facilities (e.g. coffee shops, restaurants), give directions (e.g. how to get to
9 your departure gate), update you on the status of your flight, and even check you
10 in automatically for your flight. Consider also that as you walk through the airport
11 your location changes. Your computing device, however, can receive continuous
12 location information updates so that it can continue to determine its location as
13 you move through the airport. At one point, as you pass a Starbucks coffee shop,
14 your hand held device notifies you that if you purchase a latte at Starbucks and
15 present your hand held device, you will receive a 50 cent discount on your latte.
16 In this example, the utility of the Secondary World is demonstrated. By knowing
17 where its particular customers are in its facility, SeaTac International is able to
18 provide a host of services that were not possible before.

19 Assume further that you are in the airport and your flight is canceled. You
20 must find a place to stay for the night. Accordingly, you wish to determine the
21 closest Double Tree hotel because you really like the warm chocolate chip cookies
22 they give you when you check in. You execute a search engine on your computing
23 device to find the nearest Double Tree hotel. The search engine application first
24 determines your current location in the Master World as indicated by the EID of
25 the Master World node that corresponds to your location. Executing a search, the

1 search engine application looks for a Double Tree hotel that has an attribute that
2 includes an EID that matches your EID. If it finds one, it simply indicates for you
3 the result. If it does not find one with the corresponding EID, it can use an
4 adjacent geozone to search for a Double Tree hotel. It may also provide driving
5 directions to the hotel. The search engine application was able to do this because
6 it was able to ascertain your location in the Master World. It did this quickly and
7 automatically with little or no effort from you.

8 Consider further that as you are driving from the airport to the hotel you
9 decide that you want to find the nearest Kinko's so that you can print 100 copies of
10 a presentation that you are to give in the morning. Consider that your hand-held
11 computing device is a cellular phone and that Sprint is the carrier. Sprint has
12 defined its own Secondary World that might, for example, be designated in terms
13 of cell nets. By virtue of having Sprint's Secondary World on your computing
14 device, you are able to ascertain your location in Sprint's Secondary World and,
15 accordingly, your location in the Master World. Consider that Kinko's also has a
16 Secondary World that links with the Master World. By executing a search
17 application on your device, you are able to ascertain the location of the nearest
18 Kinko's as well as driving directions thereto. All of this is possible because your
19 device has access to the Master World and one or more Secondary Worlds. In this
20 example, the Master World provides a mechanism to daisy chain two or more
21 Secondary Worlds together. This is possible because the Secondary Worlds have
22 at least one reference or link into the Master World.

Exemplary Device Architecture

Fig. 7 shows computing device 600 in somewhat more detail. In this particular embodiment, device 600 comprises an architecture that includes the following components: a location service module 602, a location provider interface 700, an application program interface (API)/Events module 702, a privacy manager 704 a location conversion module 706, one or more applications 608 and one or more location providers 606. Also included in the architecture is an active directory 708, Web service 710, location database 712, and personal places 714. The architecture can be implemented in any suitable hardware, software, firmware or combination thereof. The architecture mentioned above is advantageous in that it enables each computing device to determine its own context or location.

Common Location Provider Interface

One particularly advantageous aspect of the described embodiment is that it employs a common interface 700 that provides a standard interface through which the location providers 606 communicate. By having a common interface, the location providers are extensible (to support future providers) in that they can be dynamically added or removed from the collection of location providers. All that is required of a particular location provider 606 is that it be written to support the common interface.

In this example, there are several location providers 606. These location providers provide location information in different forms. For example, a GPS location provider might provide location information that is GPS specific. Similarly, an IP/Subnet location provider might provide information that is

1 specific to an Internet Protocol. A mobile phone location provider might provide
2 location information in the form of a cell ID. In addition, a location User
3 Interface (UI) might provide location information in the form of a user entry that
4 specifies a city, street or building. All of the location information that is provided
5 by the various location providers is processed by the location service module 602
6 so that a current device location can be determined. To determine the current
7 device location, the location service module 602 may have to consult with an
8 active directory 708, a Web service 710, or a location database 712. In the
9 illustrated example, the active directory 708 might, for example, maintain a
10 secondary world and other networking metadata such as subnet and "site"
11 information that can help determine location based on networking connectivity.
12 Web service 710 can hold the master or secondary worlds, the attributes of which
13 can be used to find location. For example, if a cell phone knows its cell tower ID,
14 then the location provider can query the secondary world to ascertain the nodes
15 that match that cell tower ID. Location database 712 is basically a version of the
16 web service that is hosted or cached locally.

17 18 **Location Providers**

19 As indicate above, the architecture contemplates multiple different location
20 providers that can provide location information to the location service module 602.
21 This information can come in many different forms and quality levels. The
22 information is then processed by the location service module 602 to determine a
23 current device location. To do this, the service module 602 ascertains from the
24 location information a particular node ID (EID and/or LUID) and a URL that is
25 associated with the tree structure with which the node is associated. Once the

1 location service module ascertains a node ID, it can then query the tree structure
2 (or more accurately a server that manages the tree structure) using the URL to
3 ascertain more information about the tree structure. For example, if the location
4 service module 602 ascertains a LUID from a particular Secondary World, it might
5 then query an active directory 708 (or an Intranet server—which is another
6 location database) to discover the parents and the children of the node. This would
7 then enable the location service module to build the Secondary World.

8 The location providers 606 can provide the location information to the
9 location service module 602 in many different ways. For example, some location
10 providers 606 may continuously provide information (e.g. the GPS provider may
11 continuously provide GPS coordinates). Alternately, the location providers can
12 periodically provide location information such as at specific times or on the
13 occurrence of definable events. For example, a user may define specific times
14 when the location information should be updated. Alternately, the location
15 information might be updated only when a device's location changes (i.e. a
16 location change event). Additionally, the location providers might provide
17 location information when polled by the location service module 602. For
18 example, the location service module 602 can call the location provider interface
19 700 and request location information from one or more of the location providers.

20 One specific location provider 606 is shown as a cache. The cache provider
21 essentially maintains a current device context or location. That is, once the
22 location service module 602 has ascertained its current location, it writes this
23 location to a cache. This enables the device 600 to ascertain its location with a
24 degree of confidence in the event all of the other location providers are not able to
25 provide location information (e.g. the GPS provider may not receive GPS

1 information because the GPS transmitter that supplies it with the information is
2 unable to contact a requisite number of satellites).

3 4 **Confidence and Accuracy Parameters**

5 One important and useful feature of the described embodiment is that one
6 or more of the location providers are configured to assign confidence parameters
7 and/or accuracy parameters to the information that they provide to the location
8 service module 602. Confidence parameters provide a measure of a provider's
9 confidence in the information that it provides to the location service module 602.
10 For example, assume that a GPS transmitter must receive information from five or
11 more satellites in order to provide highly confident information. Assume that only
12 three satellites are available at the time. The GPS transmitter would then transmit
13 its information based only on the three satellites. The GPS provider would then
14 know that the information it receives from the GPS transmitter was based only on
15 three satellites rather than the desired five or more. In this case, the GPS provider
16 can set a confidence parameter on the location information that indicates that it has
17 a lower confidence level than if the information were based on the desired five or
18 more satellites. In this case, the location service module 602 can take the
19 confidence parameters for all of the location providers into account when
20 determining the location of the device. This is discussed in more detail below.

21 With respect to the accuracy parameters, consider that the location
22 information that is received from the location providers is accurate to varying
23 degrees. Some information may be accurate to within one mile, while other
24 information may be accurate to within 100 feet. The location providers are
25 desirably configured to assign accuracy parameters to the location information that

1 they provide to the location service module 602. The accuracy parameters give
2 the location service module an indication of the accuracy of the information.

3 When the confidence and accuracy parameters are used by the location
4 service module 602, the module can make decisions on how to use the location
5 information it receives from each provider. For example, the location service
6 module 602 might disregard completely any information that has a low confidence
7 parameter. It might, on the other hand, strike a balance between the accuracy of
8 the information and its confidence. For example, the module 602 might be
9 programmed to use information with lower levels of accuracy only when there is a
10 high level of confidence in the information. The module 602 might utilize the
11 parameters to assign weights to the information so that the location is calculated as
12 a weighted function of the confidence and accuracy of the information.

13 Another use of the confidence parameters is as follows: Assume that the
14 location service module has determined a device location and has written that
15 location to a cache. At the time when the location is written to a cache, it is
16 assigned perhaps a high confidence level. Assume further that all of the other
17 location providers are unavailable to provide location information. For a period of
18 time, the location service module 602 can use the cache location as a current
19 location and be fairly confident that its information is generally accurate. In this
20 case, the location service module might assign a linearly decreasing confidence
21 level to the information over time so that at some point, it ceases to use the
22 information or informs the user that the information cannot be guaranteed.

Location, Trust, and Timestamp

When the location providers provide their information to the location service module 602, the information can include, in addition to the confidence and accuracy parameters, the actual location information in a known format, a trust parameter and a timestamp. The trust parameter is a metric that is assigned by the location service module 602 to one or more of the location providers and defines the trust that the location service module has for the particular location provider. The timestamp is a metric that defines the time when the location information was provided by the location provider. This assists the location service module 602 in ascertaining whether information is stale and might need refreshed.

Once the location service module 602 has all of the location information, it can then set about determining the location of the device.

Fig. 8 is a flow diagram that describes steps in a method of determining a device context which, in this example, is the device location. These steps are implemented by the location service module 602.

Step 800 gets the current device context. The current context can be the last calculated device context that is stored in the cache. Step 802 determines whether any of a number of context providers are available to provide context information. The location service module might do this by polling the context providers to ascertain which of the providers are active and valid. Step 804 determines whether all of the providers are inactive. If all of the providers are inactive, step 806 decreases the confidence in the current context over time and uses the current context as the device context. Step 802 then continues to monitor for current active and valid providers. If step 804 determines that one or more of the context providers are active, then step 808 orders the active and valid context

1 providers. When the location service module 602 orders or sorts the context
2 providers, it does so as a function of the confidence of the provider's information
3 and/or the trust that the location service module has in the location provider. This
4 provides a ranked list of the location providers. Step 810 checks to ascertain
5 whether the context information appears to be correct. For example, where the
6 context is the location of the device, the location service module 602 might know
7 that five seconds ago the current location was Redmond, Washington.
8 Accordingly, location information that indicates that the current location is
9 Beijing, China would be incorrect. Step 812 then determines whether any of the
10 context information conflicts with either the device's current context or the context
11 information from other providers. For example, the location service module 602
12 can compare the context information from each of the context providers with the
13 information in the cache. If any of the information conflicts with the cached
14 information, then the information from that context provider can be discarded.
15 Similarly, if context information varies inordinately as between the context
16 providers, then step 814 can select the context providers having a predefined level
17 of trust and perhaps use just their information (Step 816). If there are no conflicts,
18 then step 816 determines the current context based upon the information that is
19 provided by all of the context providers. In the described embodiment, this step is
20 implemented by using the information to map to a particular node in one or more
21 of the hierarchical tree structures mentioned above. For example, the location of
22 the device can be ascertained by mapping the information to a particular node, and
23 then completely traversing the tree structure until the root node is reached. Step
24 818 then updates the current context by perhaps writing it to the cache and returns
25 to step 802 to determine the active and valid context providers.

1 The method described above provides a way for the location service
2 module to receive location information and use only the location information that
3 appears mostly likely to represent a current location. Conflicting information can
4 be discounted or disregarded thereby assuring that only the most trusted, accurate
5 and confident information is utilized to determine the device's current location.
6

7 **Self Monitoring**

8 In addition to the confidence and accuracy parameters, one or more of the
9 location providers are advantageously programmed to self monitor their own
10 operation for various irregularities that can occur. On the occurrence of an
11 irregularity, the location providers are configured to notify the location service
12 module 602. For example, the source from which the location provider receives
13 its information may go off line for a period of time so that the location provider is
14 unable to receive any additional information. In this case, the location provider
15 might generate a "provider out" message and send it to the location service
16 module 602. When the location service module 602 receives the "provider out"
17 message, it can then take steps to exclude the location information from that
18 provider from any location calculations that it performs. When the location
19 provider's source comes back on line, it can generate a "provider on" message that
20 informs the location service module 602 that it is able to transmit location
21 information to the module. Of course, the location service module can be notified
22 by the location providers on the occurrence of other operational irregularities, with
23 the above example constituting but one specific case.
24
25

Applications

Once the location service module 602 has determined the device's location, it can receive queries from one or more applications 608. In the Fig. 7 example, the applications include a web site application, an Outlook application, and a service discovery application. In the present example, the web site application can be any web site application that is capable of rendering location-specific services. For example, the user of the device 602 might access Amazon.com's web site to buy a favorite book. When the user purchases their book, Amazon.com must now compute the taxes that the user must pay. In this example, a script executing on Amazon.com's web site might query device 602 to learn of the user's location. In this particular example, the device might respond to the query by returning the state in which the user is making the purchase. Amazon.com can then assess the tax automatically. Amazon.com might also desire to know where the individual is located so that they can select an optimal shipping method (UPS or Express Mail). Depending on where the individual is located, one method may be preferred over the other. The Outlook application might query the location service module to ascertain the location because it (or the operating system, e.g. Windows) may change device settings based on the location of the computing device. For example, the user may print on one particular printer while at work, and another particular printer when at home. When the Outlook application determines that the user has gone home for the day, it can automatically change the device settings for the printer at the user's home. It might acquire the print settings from a personal places data store 714. Thus, the device is automatically configured for use depending on the user's location. The service discovery application might query the device to determine its location so that it can render a particular service

1 depending on where the device is located. For example, if the user asks the
2 application to locate the nearest color printer, the service discovery application
3 might query the location service module to ascertain the device's current location
4 so that it can use this information and find the nearest color printer. Consider also
5 that the Outlook application could configure itself email to a work location (when
6 an individual is at work) or to a home location (when an individual is at home). In
7 addition, the Outlook calendar can become location aware, e.g. when you change
8 time zones, your appointments would show up in the proper time slots.

9 As one can imagine, the possibilities are seemingly endless. This
10 functionality is made possible through the use of the Master World and one or
11 more Secondary Worlds.

12 13 **Application Program Interface/Events**

14 In the described embodiment, the applications 608 communicate with the
15 location service module 602 through one or more application program interfaces
16 (APIs) and/or events. The applications can make function calls on the API to
17 query the location service module as to its current location. Similarly, the
18 applications can register for location notifications by using an events registration
19 process. For example, an application may register for a notification when the user
20 changes their location. Consider the case where an application requests to be
21 notified when the user arrives at work or at home so that the application can
22 change the device's configuration (such as printer configuration).

23 Fig. 9 is a flow diagram that describes steps in a method in accordance with
24 the described embodiment. The steps that are described are implemented by
25 device 600. Step 900 receives information that pertains to the current context of

1 the device. In this particular example, a portion of the information is received
2 from one or more context providers which, in this case, are location providers.
3 Step 902 processes the information on and with the device to ascertain the current
4 context of the device. In the illustrated example, the device maintains (or has
5 access to) one or more of the Master World and one or more Secondary Worlds.
6 When the device receives all of the location information, it maps the information
7 to a particular node in the hierarchical tree structure that defines the Worlds. It
8 then traverses the tree structures to ascertain the complete context (i.e. location) of
9 the device. Step 904 receives calls from one or more applications that request
10 information that pertains to the device's current context or location. In the
11 illustrated example, the applications can call one or more APIs to request the
12 information or the applications can register for event notifications. Step 906 then
13 supplies the applications with at least some information that pertains to the current
14 device location. As will be discussed below, a security policy or privacy policy
15 can be applied to the information before it is returned to the applications.

17 **Privacy Manager**

18 In one embodiment, a privacy manager 704 (Fig. 7) is provided. Although
19 the privacy manager is illustrated as being incorporated on the device, it could be
20 implemented by a trusted entity such as a trusted server that is not part of the
21 mobile computing device. The privacy manager can be implemented in any
22 suitable hardware, software, firmware or combination thereof. In the illustrated
23 example, the privacy manager comprises a software module that is incorporated in
24 the mobile computing device.

The privacy manager 704 addresses privacy concerns that are associated with the information that is collected by the computing device. Specifically, the location service module can calculate detailed information regarding the location of the computing device. It may be desirable, in some instances, to filter the information that is provided to various applications. That is, it is entirely likely that a user may not want their specific location information provided to untrusted applications. In these instances a user might just desire for location service module 602 to inform such applications that the user is in the State of Washington.

Fig. 10 shows a flow diagram that describes steps in a privacy protection method in accordance with the described embodiment. These steps can be implemented by the privacy manager 704.

Step 1000 defines a plurality of privacy levels. Exemplary privacy levels are set forth in the table immediately below:

Privacy Level	Approximate Scale	Level of Revelation
0	-	No location information is returned
10	100,000 Km	Planet/Continent
20	1,000 Km	Country
30	100 Km	State
40	10-100 Km	City & County or Region
50	10 Km	Postal Code & Phone Area Code
60	1 Km	Full Postal Code (Zip + 4) & Area Code and Exchange
70	100 m	Phone Number & Building/Floor

80	10m	Room #
90	1m	Exact Coordinates

In the illustrated table, 10 different privacy levels are defined and each has an associated approximate scale. For example, a privacy level of 0 means that no location information is returned. A privacy level of 90 means that very detailed location information is returned.

Step 1002 assigns various privacy levels to the individual nodes in one or more hierarchical tree structures. For example, each node of the Master World and the Secondary Worlds can have a privacy level associated with it. The root node of the Master World tree structure might have a privacy level of 10, while the node that represents a current location in a Secondary World might have a privacy level of 90. Step 1004 determines the context of the computing device. In the present example, the context is the device location and examples of how this is done are given above. Individual applications that call the location service module can have privacy levels associated with them. These privacy levels can be assigned by individual users. For example, a trusted application might have a privacy level of 90, while an untrusted application might have a privacy level of 30. Step 1006 receives context queries from one or more applications. Here, an application calls the location service module 602 (Fig. 7) to ascertain the location of the device. Step 1008 determines the privacy level associated with the application or applications. For example, if a untrusted application calls to request location information, the privacy manager 704 would determine that the application has a privacy level of 30. The privacy manager then traverses (step 1010) one or more hierarchical tree structures to find a node with a corresponding privacy level so

1 that it can select the information that is associated with that node. In this example,
2 the traversal might involve jumping from the Secondary World to the Master
3 World to find the node that corresponds to the state in which the user is located.
4 Once the corresponding node is found, step 1012 returns the context information
5 (e.g. location information) associated with the node. In this case, the location
6 service module would inform the application that the user's location is the State of
7 Washington.

8 As an example, consider the following: There is a web site that gives up to
9 the minute weather of various locations. Accordingly, you might assign this web
10 site a privacy level of 60 so that you can receive weather information for the
11 geographical area that corresponds to your present full postal code. Another web
12 site might be a corporation intranet web site that is a trusted web site. Thus, any
13 applications associated with this web site can be assigned a privacy level of 90 so
14 that you can give them precise location information as to your whereabouts.

15 Thus, in the present example, the computing device is able to determine the
16 source (i.e. application) of its queries and modulate the information that is returned
17 to the application as a function of the application's identity. The computing device
18 is able to do this because it has access to the Master World and one or more
19 Secondary Worlds. The above description constitutes but one exemplary way of
20 accomplishing this feat.

21 22 **Location Beacons as a Location Provider**

23 In one embodiment, one of the location providers comprises a location
24 beacon that beacons or transmits information to enable a computing device to
25 actively participate in its current context. Location beacons can comprise

1 standalone devices that can be retrofitted onto existing infrastructures, e.g. a
2 smoke detector or wall outlet in order for the device to have a power source.

3 Fig. 11 shows an exemplary beacon 1100 that is mounted on a structure
4 1102. Structure 1102 can be any suitable structure such as a wall in a conference
5 room or public place, a smoke detector, an electrical socket and the like. In the
6 described embodiment, the location beacons are small inexpensive devices that
7 can be permanently mounted in special locations such as conference rooms,
8 building lobbies, airport gates, public places and the like. The beacons announce
9 the physical location in the form of an EID and/or LUID to all mobile devices that
10 are within range, such as laptops, tablet PCs, hand held computers, mobile phones,
11 wearable computers and the like.

12 In the described embodiment, the location beacon can identify the particular
13 locations by beaconing standard information that will be understood by the mobile
14 computing devices. In the present example, the beacons can transmit one or two
15 location identifier pairs comprising an EID/URL pair and a LUID/URL pair. The
16 beacon might also transmit multiple LUIDs. The EID and LUID give the present
17 node location in the Master World and Secondary World respectively. The URLs
18 provide a reachable location for the Master and Secondary Worlds. For example,
19 the URL associated with the Secondary World can give a service location that the
20 device can use to query information about the Secondary World so that it can
21 derive its context and take advantage of resources or services that are associated
22 with the nodes in the Secondary World.

23 The beacons can also transmit a digital signature that can be used by the
24 device to ascertain that the beacon is valid and legitimate. Any suitable signature
25 or verification method could be used. In addition, and of particular use in the

1 context-aware environment, the beacon can be programmed to transmit code
2 download pointers to devices within range. The code download pointers can
3 enable the computing device to access software code that permits them to interact
4 with their environment. Consider the following example: You walk into a
5 conference room with your cell phone computing device and immediately a
6 beacon in the conference room transmits your location in the form of an EID/URL
7 pair and a LUID/URL pair. Your device uses the information pairs to ascertain its
8 location in the Master and Secondary Worlds as described above. The beacon also
9 transmits a code download pointer that points to software code that enables you to
10 operate the video projector in the conference room using your hand-held cellular
11 phone. In this manner, the beacon serves as more than just a location beacon—it
12 permits you, through your computing device, to actively participate in your
13 surroundings.

14 The beacons can transmit the information in any suitable way, e.g. wireless
15 methods including infrared and radio frequencies. In one embodiment, Bluetooth
16 short range radio frequency communication can be used to provide a low cost, low
17 power alternative.

18 **Context-Aware Vehicle**

19
20 Fig. 12 shows an exemplary vehicle (e.g., automobile) computer system
21 1200 that can be used in connection with the context-aware systems and methods
22 described above and below. System 1200 is configured for use in a vehicle such
23 as a car, truck or other similar conveyance. When in place in a vehicle, computer
24 system 1200 can provide a multitude of context-aware services as will become
25 apparent below.

1 Vehicle computer system 1200 has a centralized computer 1202 coupled to
2 various external peripheral devices, including a display device 1204, security
3 sensors 1206, a vehicle diagnostic interface 1208, speakers 1210, a vehicle battery
4 1212, a backup battery 1214, and antenna(s) 1216. The computer 1202 is
5 assembled in a housing 1218 that is sized to be mounted in a vehicle dashboard,
6 similar to a conventional automobile stereo. In the illustrated example, the
7 housing 1218 has a form factor of a single DIN (Deutsche Industry Normen).
8 Alternatively, it could be housed in a 2 DIN unit or other special form factor for an
9 OEM.

10 The computer 1202 runs an open platform operating system which supports
11 multiple applications. Using an open platform operating system and an open
12 computer system architecture, various software applications and hardware
13 peripherals can be produced by independent vendors and subsequently installed by
14 the vehicle user after purchase of the vehicle. This is advantageous in that the
15 software applications do not need to be specially configured for uniquely designed
16 embedded systems. In the illustrated example the open hardware architecture runs
17 a multitasking operating system that employs a graphical user interface. A
18 multitasking operating system allows simultaneous execution of multiple
19 applications. One such operating system is the "Windows" brand of operating
20 systems (e.g., the "Windows CE" operating system) sold by Microsoft
21 Corporation of Redmond, Washington.

22 The computer 1202 can include at least one storage drive which permits the
23 vehicle user to download programs and data from a storage medium. In the
24 illustrated implementation, the computer 1202 has a CD ROM (or other optical)
25 drive 1220 which reads application-related CDs, as well as musical, video, game,

1 or other types of entertainment CDs. The computer 1202 may also include other
2 storage devices, such as a magnetic disk drive, smart card reader, PCMCIA card
3 sockets, a hard disk drive, or a DVD ("digital video disk" or "digital versatile
4 disk") drive.

5 The storage drives are mounted in a base unit 1228 of the housing 1218.
6 The base unit 1228 is constructed and sized to be mounted in the dashboard.
7 Optionally, this base unit may be removable in the same fashion as a laptop
8 computer and its associated docking station. This option allows the user to take
9 the vehicle computer to his/her home or office to serve as his/her portable PC.
10 The housing 1218 also has a faceplate 1230 which is pivotally mounted to the
11 front of the base unit 1228 and may optionally be detachable. The faceplate can
12 be rotated to permit easy and convenient access to the storage drives.

13 The computer 1202 has a keypad 1232 and a display 1234 on the faceplate
14 1230. The operating system executing on the computer 1202 controls the
15 faceplate peripheral, which, through the faceplate processor, can control the
16 faceplate keys 1232 and the faceplate display 1234 as peripheral devices when the
17 faceplate is attached to the base unit. Additionally, the computer 1202 has a voice
18 recognition device to permit the user to verbally enter commands in a hands-free,
19 eyes-free environment. These voice commands can be used for controlling most
20 operating modes of the vehicle computing platform. The computer 1202 is also
21 equipped with an IrDA (infrared developers association) transceiver port 1236
22 mounted on the faceplate 1230 to transmit and receive data and programs using
23 infrared signals. The entire faceplate unit 1230 behaves as a multifunction
24 peripheral to the computing platform.
25

1 The computer 1202 can output visual data to the LCD 1234 at the faceplate,
2 or to the display device 1204. In the exemplary illustration, display 1234 is a back
3 lit LCD and display 1204 is a small flat panel display (e.g., 6.4" screen) that is
4 movably mounted on a stand or yoke and remotely located from the computer.
5 Additional display devices may also be added that are similar to display 1204 or
6 1234. Different types of display devices may also be added, such as a Heads Up
7 Display (HUD).

8 The display 1204 is fully adjustable to different viewing positions that can
9 be seen by the driver or other passengers in the vehicle. The type of data
10 displayed can range widely from word instructions concerning the vehicle's
11 performance, to diagrammatic directions from a navigation system, to video
12 movies for in-car entertainment. The display 1204 can be equipped with an
13 automatic override switch 1238 which automatically disables the display of any
14 non-driving related data when positioned to be viewed by the driver. When facing
15 the driver, only information supportive and helpful to driving (e.g., diagnostics,
16 navigation directions) is displayed on the monitor, while distracting information
17 (e.g., video movies, games) is blocked from display. In one implementation, the
18 switch is an electrical cylindrical switch which closes when the display is capable
19 of being viewed by the driver; thus, the software can sense the display position
20 and only allow permitted information to be displayed.

21 In general, the vehicle computer system 1200 can be used to integrate
22 multiple vehicle-related systems onto one open platform hardware and software
23 architecture. For instance, the vehicle computer system 1200 can serve as a
24 multimedia entertainment system, a navigation system, a communications system,
25 a security system, and a diagnostics system. Moreover, the vehicle computer

1 system 1200 provides additional functionality traditionally associated with desk-
2 top and laptop personal computers. For instance, vehicle computer system 1200
3 can support word processing applications, spreadsheet applications, database
4 applications, web browser applications, and appointment/schedule applications.
5 Furthermore, the vehicle computer system 1200 can be configured to operate as a
6 server to other computing units in the vehicle to distribute games, video movies,
7 and the like to passengers. Furthermore, the system includes an internet
8 connectivity engine 1240 that is configured to establish connectivity with the
9 Internet. This can be done in any suitable way. For example, the internet
10 connectivity engine 1240 can wirelessly establish connectivity with the Internet
11 through known wireless techniques.

12 In accordance with the described embodiment, information can be
13 displayed on either display device 1204 or display 1234. The information can be
14 provided by an application running on computer 1202, or by a device external to
15 computer 1202, such as sensors 1206 or via diagnostic interface 1208, antenna
16 1216, IrDA port 1236, etc. Such information can also be provided to the computer
17 via the internet connectivity engine 1240. Information that can be displayed
18 includes any type of data or control information. Additionally, information to be
19 displayed can include a "caption" or "label" that describes the data. Examples of
20 data that can be displayed include street addresses, phone numbers, and directions
21 (e.g., "Turn Left At Light On Main Street"). Such data can be displayed either
22 including a caption describing the data (e.g., "Address: 12345 Washington
23 Street", where "Address:" is the caption portion of the information) or without a
24 caption (e.g., "12345 Washington Street"). Examples of control information
25

1 include toolbars, menu options, and user-selectable on-screen regions (such as
2 buttons), as well as instructions, headings, and other descriptive information.

3 In the discussion herein, the various embodiments are described in the
4 general context of computer-executable instructions, such as program modules,
5 being executed by one or more conventional personal computers. Generally,
6 program modules include routines, programs, objects, components, data structures,
7 etc. that perform particular tasks or implement particular abstract data types.
8 Moreover, those skilled in the art will appreciate that the described embodiments
9 can be practiced with other computer system configurations, including hand-held
10 devices, multiprocessor systems, microprocessor-based or programmable
11 consumer electronics, network PCs, minicomputers, mainframe computers, and
12 the like. In a distributed computer environment, program modules may be located
13 in both local and remote memory storage devices.

14 Fig. 13 shows exemplary components of computer 1202 of Fig. 12 in more
15 detail. Computer 1202 includes one or more processors or processing units 1352,
16 a system memory 1354, and a bus 1356 that couples various system components
17 including the system memory 1354 to processors 1352.

18 The bus 1356 represents one or more of any of several types of bus
19 structures, including a memory bus or memory controller, a peripheral bus, an
20 accelerated graphics port, and a processor or local bus using any of a variety of
21 bus architectures. The system memory 1354 includes read only memory (ROM)
22 1358 and random access memory (RAM) 1360. A portion of the operating system,
23 such as kernel 1362, contains the basic routines that help to transfer information
24 between elements within computer 1202, such as during start-up, is stored in ROM
25 1358.

1 A number of program modules can be stored in ROM 1358 or RAM 1360,
2 including an operating system 1364 and one or more application programs 1366.
3 A user may enter commands and information into computer 1202 through various
4 input devices, such as a keyboard (e.g., keypad 1232 of Fig. 12), touchscreen,
5 pointing device, microphone, joystick, game pad, satellite dish, scanner, or the like
6 (not shown in Fig. 13). These and other input devices are coupled to the
7 processing unit 1352 through an input/output (I/O) interface 1368 that is coupled
8 to the bus 1356. A display 1204 or 1234, or other type of display device, is also
9 connected to the bus 1356 via an interface, such as a video adapter(s) 1370. Data
10 to be displayed on display 1204 or 1234 is provided to adapter 1370 by a display
11 generator 1372 of operating system 1364. In addition to the display, computers
12 can include other peripheral output devices (not shown in Fig. 13) such as
13 speakers and printers that are coupled to the processing unit 1352 through I/O
14 interface 1368.

15 Generally, the processors of computer 1202 are programmed by means of
16 instructions stored at different times in the various computer-readable storage
17 media of the computer. Programs and operating systems are typically distributed,
18 for example, on floppy disks or CD-ROMs. From there, they are installed or
19 loaded into the secondary memory of a computer. At execution, they are loaded at
20 least partially into the computer's primary electronic memory. For example, the
21 above-described context and location service can be stored in this manner so as to
22 permit the computer to make full use of the context-aware (location-aware)
23 applications described above and below.

24 The embodiments described herein include these and other various types of
25 computer-readable storage media when such media contain instructions or

1 programs for implementing the steps described below in conjunction with a
2 microprocessor or other data processor. The described embodiments also include
3 the computer itself when programmed according to the methods and techniques
4 described below. Furthermore, certain sub-components of the computer may be
5 programmed to perform the functions and steps described below. The described
6 embodiments include such sub-components when they are programmed as
7 described. In addition, the described embodiments include data structures,
8 described below, as embodied on various types of memory media.

9 For purposes of illustration, programs and other executable program
10 components such as the operating system are illustrated herein as discrete blocks,
11 although it is recognized that such programs and components reside at various
12 times in different storage components of the computer, and are executed by the
13 data processor(s) of the computer.

Exemplary Context-based Architecture

Fig. 14 shows an exemplary architecture 1400 that can be utilized to implement a context-aware vehicle. Some of the components of the architecture are similar to or the same as components that have already been discussed above. To this extent, these components are not described in additional detail here. These components include an application program interface/events component 1402, context service module 1404, context provider interface 1406, and privacy manager 1408.

Additional components that are germane to this discussion include one or more applications 1410 that are executable on the vehicle's computer, one or more context providers 1412, a context data store 1414, and a behavior engine 1418.

Applications 1410 can include any suitable applications that are capable of being executed in connection with a vehicle such as an automobile. Exemplary applications can include those that are used to implement or be used in connection with audio/visual equipment or peripheral devices such as the radio, entertainment devices, visual displays, temperature regulation devices, and various controllers that can be associated with equipment on the vehicle, to name just a few.

The context providers 1412 can be any suitable context providers that are capable of providing context information, via the context provider interface 1406 and context service module 1404, to the behavior engine 1418. Context providers can include such things as environmental sensors (e.g. light, heat, smog and the like). It will be appreciated that although the context providers are shown as being on the computing device itself (i.e. located on the vehicle), the context providers

1 can be located off of the vehicle. For example, context providers can broadcast
2 wireless context information for receipt by the computing device.

3 Context store 1414 can store per-user context preferences that are available
4 locally and/or it can cache per-user context preferences that are present somewhere
5 outside the vehicle. This store can also be used to store context preference
6 information that is common to all users (per-vehicle). This store can also be used
7 to temporarily store all the context data that is being provided by the various
8 context providers. This way, the behavior engine can act on this data at a later
9 time, rather than instantaneously. In addition, the store can hold per-vehicle data.
10 User-specific context information can pertain to any suitable type of context
11 information. As an example, the user might provide their preferences for seat
12 settings, vehicle temperature, vehicle interior lighting, radio volume and the like.
13 These preferences can be maintained on the vehicle in a vehicle store, or
14 externally of the vehicle in, for example, a network accessible data store (e.g. a
15 data store accessible via a Web server).

16 Behavior engine 1418 comprises a software module that receives context
17 information from one or more sources and manipulates in some way or causes the
18 behavior of the vehicle to be manipulated. Manipulation of the vehicle's behavior
19 can manifest itself, in the illustrated example, through settings or other parameters
20 that are passed from the behavior engine 1418 to one or more applications 1410
21 and that affect the behavior of the vehicle. The settings or parameters thus cause
22 the application to function, in some way, in a manner that is unique to the context
23 in which the vehicle is determined to be.

24 As an example consider the following: A vehicle is traveling on a long trip
25 that started early in the morning. As dusk approaches, a context provider 1412 in

1 the form of a light sensor provides information to the context service module 1404
2 which, in turn, ascertains that outside lighting conditions are sufficient to require
3 the headlights to be turned on. Accordingly, the behavior engine 1418 causes the
4 application associated with the headlights to turn the headlights on. In this
5 specific example, the behavior engine can also cause the application associated
6 with the driver's display to change the light contrast setting so that the driver can
7 more easily view the display. This solution, as well as the other solutions
8 described below, can be advantageously implemented in software so that the
9 necessity of building any hardwired hardware solutions is eliminated.

10 As another example, consider the following: While traveling in a context-
11 aware vehicle, the vehicle enters an area where there is a large amount of smog. A
12 context provider in the form of a smog sensor provides information to the context
13 service modules which, in turn, determines that smoggy conditions exist. As a
14 result, the behavior engine 1418 causes the application 1410 associated with the
15 vehicle's ventilation system to automatically shut off the outside supply of air.

16 Fig. 15 is a flow diagram that describes steps in a method in accordance
17 with the described embodiment. The method can be implemented in any suitable
18 hardware, software, firmware, or combination thereof. In the illustrated example,
19 the method is implemented in software. In a specific implementation, this method
20 can be implemented in connection with the architecture shown in Fig. 14.

21 Step 1500 receives context information. This information can comprise any
22 suitable type of information and can be received in any suitable way. For
23 example, the context information can comprise user-specific information. This
24 user-specific information can be received from a user who physically enters the
25 information into the vehicle's computer. Alternately, the user-specific information

1 can be retrieved from a data store, either onboard or remote from the vehicle. If
2 the information is onboard the vehicle, the information can simply be retrieved
3 from whatever data store holds the information. If the information is maintained
4 remote from the vehicle (e.g. in a Web-accessible data store), the information can
5 then be retrieved in any suitable way such as via a wireless connection with the
6 Web that is then used to access the remote data store and retrieve the information.

7 The context information can also comprise information associated with the
8 vehicle's operating environment and can be gleaned from various different sensors
9 that can be provided on the vehicle. Examples of this are given above.

10 Step 1502 determines, from the context information, a current vehicle
11 context. This step can be implemented by a context service module such as
12 module 1404 in Fig. 14. The current context can be determined by ascertaining
13 the current vehicle user. The current context can also be determined by making a
14 decision, based upon the context information, about the likely context of the
15 vehicle. For example, if the context information suggests that it is likely dark
16 outside of the vehicle, a decision concerning the current context can be made with
17 an ascertainable degree of certainty. Based on the current vehicle context, the
18 behavior of the vehicle can then be modified in some way or manner that is
19 consistent with the vehicle's current context (step 1504). This step can be
20 implemented by a behavior engine such as engine 1418 in Fig. 14. The behavior
21 engine receives information from the context service module 1404 indicating the
22 current vehicle context. Based on this, the behavior engine can cause settings or
23 other parameters to be passed to applications associated with the context so that
24 the applications can modify the behavior of the vehicle in some way.

Exemplary Location-Based Architecture

Fig. 16 shows an exemplary architecture 1600 that can be utilized to implement a location-aware vehicle. Some of the components of the architecture are similar to or the same as components that have already been discussed above. To this extent, these components are not described in additional detail here. These components include an application program interface/events component 1602, location service module 1604, location provider interface 1606, and privacy manager 1608.

Additional components that are germane to this discussion include one or more applications 1610 that are executable on the vehicle's computer or on the vehicle itself, one or more location providers 1612, a context store 1614, and a behavior engine 1618.

Applications 1610 can include any suitable applications that are capable of being executed in connection with a vehicle such as an automobile. Exemplary applications can include those that are used to implement or be used in connection with a location-sensitive service. For example, one application can comprise an Internet connectivity application that establishes an Internet connection via a local ISP. For example, the Internet connectivity engine 1240 (Fig. 12) can call a toll free telephone number to get a list of affiliated ISPs and a set of <location, phone number> pairs. This list can also be pre-installed in the vehicle's computer, or downloaded periodically via that Internet. When the user tries to access the Internet, the Internet connectivity engine 1240 calls the location service module 1604 to ascertain the current location. Based on the current location, the Internet connectivity engine 1240 selects an access number that is local to the area in which the vehicle is located. As the vehicle moves between locations, the Internet

1 connectivity engine 1240 can get location updates and update its local access
2 number. Another application can control the radio and its various settings, as will
3 be discussed in more detail below in the section entitled "Radio Station Settings".

4 The location providers 1612 can be any suitable location providers that are
5 capable of providing location information, via the location provider interface 1606
6 and location service module 1604, to the behavior engine 1618. Exemplary
7 location providers are discussed above. In vehicle-based applications, however,
8 one source of location information comprises a GPS source that provides GPS
9 information to an onboard location provider.

10 Context store 1614 can store per-user location preferences that are available
11 locally and/or it can cache per-user location preferences that are present
12 somewhere outside the vehicle. This store can also be used to store location
13 preference information that is common to all users (per-vehicle). This store can
14 also be used to temporarily store all the location data that is being provided by the
15 various location providers. This way, the behavior engine can act on this data at a
16 later time, rather than instantaneously. In accordance with the described
17 embodiment, users can provide user-specific location information that is
18 maintained in store 1614. This user-specific location information can pertain to
19 any suitable type of location information. As an example, and one that is explored
20 in more detail below, a user might enter their favorite radio stations for several
21 different cities that they travel to often. These preferences can be maintained and,
22 when the location service determines the location of the user's vehicle to be one of
23 the subject cities, the behavior engine 1618 causes the application that is
24 associated with the station controller to automatically program the radio stations
25 for the new city.

1 Behavioral engine 1618 comprises a software module that receives location
2 information from one or more sources and manipulates in some way the behavior
3 of the vehicle. Manipulation of the vehicle's behavior can manifest itself, in the
4 illustrated example, through settings or other parameters that are passed from the
5 behavior engine to one or more applications 1610. The settings or parameters thus
6 cause the application to function, in some way, in a manner that is unique to the
7 location in which the vehicle is determined to be.

8 Fig. 17 is a flow diagram that describes steps in a method in accordance
9 with the described embodiment. The method can be implemented in any suitable
10 hardware, software, firmware, or combination thereof. In the illustrated example,
11 the method is implemented in software. In a specific implementation, this method
12 can be implemented in connection with the architecture shown in Fig. 16.

13 Step 1700 receives location information. This information can comprise
14 any suitable type of information and can be received in any suitable way. For
15 example, the location information can comprise user-specific information. This
16 user-specific information can be received from a user who physically enters the
17 information. Alternately, the user-specific information can be retrieved from a
18 data store, either onboard or remote from the vehicle. If the information is
19 onboard the vehicle, the information can simply be retrieved from whatever data
20 store holds the information. If the information is maintained remote from the
21 vehicle (e.g. in a Web-accessible data store), the information can then be retrieved
22 in any suitable way such as via a wireless connection with the Web.

23 Step 1702 determines, from the location information, a current vehicle
24 location. This step can be implemented by a location service module such as
25 module 1604 in Fig. 16. Based on the current vehicle location, the behavior of the

1 vehicle is modified in some way (step 1704). This step can be implemented by a
2 behavior engine such as engine 1618 in Fig. 16. The behavior engine receives
3 information from the location service module 1604 indicating the current vehicle
4 location. Based on this, the behavior engine can cause settings or other parameters
5 to be passed to applications associated with the location so that the applications
6 can modify the behavior of the vehicle in some way.

8 **Radio Station Settings**

9 In one embodiment, a user can specify radio stations and locations
10 associated with those radio stations. Based on the location of the vehicle, the user-
11 specified radio stations for the particular radio stations are mapped automatically
12 to the buttons of the radio when the vehicle's location coincides with a location
13 specified by the user. As an example, consider the following: Assume that 950
14 AM is the all-sports radio channel in Seattle and 880 AM is the all-sports radio
15 channel in Portland. Further assume that the user has specified that button #1 of
16 the radio is to be associated with 950 AM when his vehicle is located in Seattle
17 and 880 AM when his vehicle is located in Portland. These user specifications are
18 referred to as the user's preferences. The user's preferences are stored in an
19 accessible location such as a per-user location store 1614. Assume now that the
20 user has to travel from Portland to Seattle on business. When the location service
21 on board the user's vehicle ascertains that the vehicle's location is Seattle, the
22 behavior engine 1618 causes the application associated with the radio station
23 controller to map Seattle station 950 to button #1 of the radio. This happens
24 automatically without the need for any user intervention.

Fig. 18 is a flow diagram that describes steps in a method in accordance with one described embodiment. The method can be implemented in any suitable hardware, software, firmware, or combination thereof. In the described embodiment, the method is implemented in software. Suitable software is described above in connection with Figs. 16 and 17.

Step 1800 associates one or more radio stations with one or more locations. This step can be implemented by a user establishing such an association by inputting their preferences into the vehicle's computer. This step can also be implemented by a user storing their preferences in some other location that is accessible by the vehicle's computer (e.g. a Web-accessible location). The preferences are then stored in a per-user location store such as the one described above. The store can either be onboard a vehicle or remote from the vehicle. If the store is remote from the vehicle, then the vehicle's computer can establish a suitable connection, for example via the Internet, and fetch the user's preferences.

Step 1802 determines the location of the vehicle. Any suitable method or technique can be used to determine the vehicle's location. In this specific example, a location service, implemented in software, can be used. Recall that the location service discussed in this document makes use of hierarchical tree structures in the form of Primary and Secondary Worlds that permit the location service to ascertain its location based upon input from one or more location providers. Once the location is determined, radio stations that are associated with that location are mapped (step 1804) to radio station buttons on the vehicle's radio. This step is implemented, in the present example, through the interaction between the behavior engine 1618 and an appropriate application 1610 that manages the mapping of the radio stations. Once the mapping has taken place, the user is free

1 to select from among the station that they previously specified. Step 1806 then
2 determines whether the vehicle is at a new location. If so, the method determines
3 the new location through the use of the location service and then branches to step
4 1804 to map any new stations associated with the new location to the radio's
5 buttons. If there is no new location, the method does nothing and simply returns
6 to step 1806.

7 The embodiment described immediately above provides a system and
8 methodology by which a user can make radio station preferences for different
9 locations known to the user. These known preferences are then enforced
10 automatically as a function of a vehicle's location.

11 In certain circumstances, however, individuals may be unaware of the
12 specific radio stations in all of the potential locations to which they may travel that
13 play the *type* of music or content to which they desire to listen. For example, an
14 individual may not know the specific radio stations in Fresno, San Francisco,
15 Sacramento, and San Diego that play the type of music that they prefer.
16 Accordingly, one embodiment allows the user to associate each of the radio
17 buttons with a radio station type. Exemplary radio station types include, without
18 limitation, those that play light rock, heavy metal, all-news, all-sports, classical,
19 talk shows and the like. The vehicle's computer maintains or has access to a list
20 that associates radio station types, location, and radio station frequency (i.e. station
21 number). This list can be resident on the vehicle's computer, or can be
22 periodically downloaded from the Internet. Now, when a user travels from
23 location to location, the location service determines the vehicle's location as
24 described above, and notifies the vehicle's behavior engine. The behavior engine
25 then causes the appropriate application that manages the radio station selections to

1 map the new radio station frequencies (or settings) to appropriate buttons on the
2 vehicle's radio. In this way, a user need simply specify radio station types and
3 need not know specific radio station settings.

4 Fig. 19 is a flow diagram that describes steps in a method in accordance
5 with one described embodiment. The method can be implemented in any suitable
6 hardware, software, firmware, or combination thereof. In the described
7 embodiment, the method is implemented in software. Suitable software is
8 described above in connection with Figs. 16 and 17.

9 Step 1900 associates at least one radio station type with a radio button.
10 This step can be implemented by a user specifying their associations or
11 preferences. For example, hard rock can be associated with button #1, soft rock
12 with button #2, and so on. These associations or preferences are then either
13 maintained by or accessible to the vehicle's computer. Step 1902 determines the
14 vehicle's location. This step can be implemented as described above. Step 1904
15 maps, for a given location, each radio station frequency for a radio station type to
16 a radio station button. So, for example, as a user drives from San Francisco to San
17 Diego, their preferences for radio station types (i.e. classical, rock, sports and the
18 like) are automatically and seamlessly made available to the user via the radio
19 buttons.

21 **Location-based Information**

22 In another embodiment, the location of a vehicle is utilized as the basis to
23 provide the driver with information that can be of interest to them. Specifically,
24 the location service onboard the vehicle determines the vehicle's location as set
25 forth above. Once the location of the vehicle is ascertained, various information

1 associated with or pertaining to that location or locations nearby can be presented
2 to the driver. This information can be any information that might be of interest to
3 the driver. For example, such information can include, without limitation,
4 upcoming weather, traffic reports, rest stops, suggestions regarding food, gas
5 stations, scenic viewpoints and the like. As an example, consider the following: A
6 user indicates in their preferences that they prefer to fill their car with gas from
7 Texaco stations only. While traveling on a trip, the location service determines the
8 vehicle's location. In conjunction with determining the vehicle's location, an
9 application executing on the vehicle's computer informs the driver that there are 3
10 Texaco stations within the next 100 miles, and the distance to each of the stations.
11 The location service can ascertain this through the use of Primary and Secondary
12 World structures as well as IDs that are associated with each of the Texaco
13 stations, as described above in detail. It will be appreciated that the information
14 that is provided to the user is different from and, in some cases independent of
15 known navigation information (i.e. street maps and directional information and the
16 like). The information that is provided to the users can be associated with one or
17 more user-specified preferences. For example, the user might have specified that
18 they only like to eat at Denny's, Stuckeys, and McDonald's while on the road.
19 Accordingly, once the location service determines the location of the user's
20 vehicle, the user can be advised of the nearest Denny's, Stuckeys, and
21 McDonald's.

22 Fig. 20 is a flow diagram that describes steps in a method in accordance
23 with one described embodiment. The method can be implemented in any suitable
24 hardware, software, firmware, or combination thereof. In the described
25 embodiment, the method is implemented in software.

1 Step 2000 determines a vehicle location. Location can be determined as
2 described above. Step 2002 provides location-specific information for the user or
3 driver of the vehicle. This information can be provided in any suitable way. In
4 one particularly useful implementation, the location-specific information can be
5 provided via a display that the user can view. The location-specific information
6 can be information that is generally provided to all drivers (i.e. upcoming
7 viewpoints, rest stops etc.). Alternately, the information can be user-specific and
8 can be based on user-provided preferences.

10 User as a Context

11 In one embodiment, one context of a vehicle comprises the user or driver,
12 as was mentioned above. That is, each user or driver is presumed to have
13 preferences that are associated with a user-specific experience that they desire
14 when they use the vehicle. For example, one user might have desired seat settings,
15 while another driver might have different seat settings. When the context is
16 location, different drivers might have different preferences that indicate where
17 they desire to eat, points of interest, radio station types, and the like.

18 Consider the case when the vehicle is a rental car. In this instance, the
19 vehicle is subject to use by a large number of users. Each user most assuredly will
20 have different preferences that govern their user experience.

21 Fig. 21 is a flow diagram that describes steps in a method in accordance
22 with one described embodiment. The method can be implemented in any suitable
23 hardware, software, firmware, or combination thereof. In the described
24 embodiment, the method is implemented in software.

1 In accordance with one embodiment, user preferences are collected (step
2 2100) and stored (step 2102) or otherwise maintained in a vehicle-accessible
3 location. For example, if the vehicle is a family car, the preferences of family
4 drivers might be stored in a locally maintained store. If the vehicle is a rental car,
5 the preferences might be stored in a large preference store remote from the
6 vehicle. One example of a remote store is one that is accessible via the Internet
7 through, for example, a Web server. For example, Fig. 22 shows a system
8 generally at 2200 that comprises multiple vehicles 2202. Each vehicle is equipped
9 with a vehicle computer 2204 such as that described above. A central server 2206
10 is accessible via the Internet and has access to a data store 2208 in which user
11 preferences are stored. Each vehicle's computer can wirelessly access the Internet,
12 as described above, so that the computer can gain access to various user
13 preferences.

14 In this embodiment, a user of the vehicle first authenticates or identifies
15 themselves to the vehicle (step 2104) at which point their preferences are retrieved
16 (step 2106). As noted above, the preferences can be retrieved from a local storage
17 location, or from a remote storage location. A user can authenticate themselves in
18 any suitable way. For example, a password might be used. Alternately, an
19 authentication service like Microsoft's Passport service might be used which
20 authenticates the user via the Internet. Further, alternately, authentication can be
21 implemented through fingerprinting techniques, voice recognition, or any other
22 suitable biometrics techniques. Consider also the following example. A user is a
23 rental car customer and is issued a "Preferred" card. This preferred card uniquely
24 identifies the user. When the user rents a car, the attendant at the rental car counter
25 swipes their card through a card reader which identifies the customer. Based on

1 the customer's identity, a computer retrieves the customer's preferences from a
2 store and communicates the preferences to a car that is waiting for the customer.

3 Once the user is authenticated and the user's preferences are retrieved, the
4 behavior engine (1418 or 1618) causes applications executing on the vehicle's
5 computer and which are associated with the subject of the user's preference, to
6 provide a user-specific experience (step 2108). This user-specific experience can
7 involve everything from adjusting the seat settings, to automatically mapping
8 radio station types to radio buttons, to providing location-specific information that
9 may likely be of interest to the user. In the rental car example above, when the
10 customer arrives at their car, their preferences have already been implemented and
11 a user-friendly car is waiting.

12 Accordingly, by maintaining or having access to user-specified preferences,
13 the user's vehicle experience can be automatically tailored to a user's likes, while
14 avoiding a user's dislikes. User convenience is enhanced in that once their
15 preferences are established, they need not worry about configuring a vehicle to
16 suit their preferences. This can all be done automatically for them.

17 18 **Conclusion**

19 The embodiments described above provide a uniform, standardized way to
20 enhance the world of context aware computing. The embodiments provide a way
21 for individuals to uniquely experience the world around them by ascertaining their
22 location in the world in a standard way. The embodiments also provide a way for
23 service providers to uniquely position their goods and services in a manner that is
24 sensitive to and appreciates the contexts, e.g. locations, of various consumers of
25 the goods and services. Unique and useful architectures and data structures are

1 employed to facilitate the user's computing experience and provide for an
2 individual-centric experience.

3 Although the invention has been described in language specific to structural
4 features and/or methodological steps, it is to be understood that the invention
5 defined in the appended claims is not necessarily limited to the specific features or
6 steps described. Rather, the specific features and steps are disclosed as preferred
7 forms of implementing the claimed invention.